Interactive Immunity and Adaptive Innovation

Luis O. Arata
Department of Fine Arts, Languages & Philosophy
Quinnipiac University, Hamden, CT 06518
E-mail: luis.arata@quinnipiac.edu

ABSTRACT

Autonomous systems must have the ability to operate on their own in dynamic, uncertain environments without breaking down. This paper presents concepts for the design and evaluation of self-repairing systems that draw on current work on immune mechanisms and artificial immune systems.

To survive in a dynamic environment we say that a constructed system must be able to adapt. The problem is how to implement the adaptive drive. This paper looks into one possible route: to model and implement features inspired by the immune system as a problem solving mechanism. This route does not exclude other adaptive mechanisms and can complement them.

The autonomous system detects malfunctions and tries to fix them on its own, tinkering with all it has at hand. Internal innovation happens when a new type of malfunction is fixed. An interesting aspect of this process is that what causes a new type of malfunction can be the result of the system's interaction with a new environment. Therefore this immune response mechanism functions as an adaptive drive. The system remembers the solution for future use and with quicker response. It has innovated with respect to its previous capabilities, and learns from this action. In this sense the system has adapted.

KEYWORDS: Interaction, innovation, immunity, imperfection, adaptation, evolution, repair mechanisms.

1. INTRODUCTION

A hurdle in the construction of an autonomous system is to find and implement design that would allow them to function in changing environments. The system has to operate according to its tasks and at the same time deal with unexpected situations that happen within its territory, for which it has no preset way to find solutions. If the system begins to malfunction while interacting with a new situation, it should be able to recognize that there is a problem, be involved in diagnosing the problem, and try to implement a solution.

Beyond its ability to carry out tasks, an autonomous system has to change in response to learning from unexpected situations. The system has to innovate at least with respect to itself. It has to be adaptable. This is one measure of its intelligence given the environment in which it is meant to perform.

Our own natural intelligent systems are a source of inspiration for designing autonomy. We appear to have two rather distinct systems tailored for different sorts of tasks. Our brain/nervous system is more for navigation and exploration. Llinás described our feeling of self as a way to help us have a sense of location in the maps of the world we create. This self we create helps us move as safely as possible in our environment [1]. Perhaps the most critical function of our brain is to help manage our body's interactions with the world.

Our immune system's task is more internal. It has to keep the body working well as it moves and interacts in its changing environment. Its function is mostly to manage internal interactions. Not surprisingly, its architecture is quite different from the nervous system. The architectures of these two intelligent systems are based on their tasks. We use both at the same time in a loosely coordinated way.

In this paper I will focus on intelligent features of our immune system that may be applied or emulated in the design of constructed autonomy. Immune mechanisms can be a vast source of inspiration for the design of autonomous systems for dynamic environments. Our natural immune system is an intelligent system in its own right, along with our brain/nervous system. It uses special procedures of its own that we can try to imitate to some extent in the design of artificial intelligence features. As we examine these features, it helps to keep in mind that they are not meant to replace existing architectures. What they could do is increase the robustness of current designs and therefore enable them to be more interactive and useful. We also have to realize that the drawback of such features is that they increase the complexity and cost of the system. The ultimate implementation of immune features would then depend on a balanced approach to the task at hand. This balance includes resources available and cost.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **AUG 2004**	2. REPORT TYPE	3. DATES COVERED **00-00-2004 to 00-00-2004**

4. TITLE AND SUBTITLE **Interactive Immunity and Adaptive Innovation**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Quinnipiac University,Department of Fine Arts, Languages & Philosophy,Hamden,CT,06518**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
Proceedings of the 2004 Performance Metrics for Intelligent Systems Workshop (PerMIS ?04), Gaithersburg, MD on August 24-26 2004

14. ABSTRACT
see report

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**6**	

2. MODELS OF IMMUNITY

Alexander Tarkanov has done extensive work in mathematical modeling of immunocomputing. To help conceptualize this field, he first groups the various approaches to the construction of intelligent systems into two directions. One uses generative grammars and logic. The other is based on biological models. In this second approach there are three main directions of research: neural networks, cellular automata, and genetic algorithms [18]. Tarkanov indicates that his work belongs to a fourth direction that is starting to gain momentum: artificial immune systems (AIS), which includes immunocomputing [2].

Classifications come in all shapes and forms. This one helps locate AIS within the many areas of AI. In this paper I will review selected features of our immune system implemented in autonomous systems that operate in changing environments. I will argue that these features of the immune system help it adapt when confronted with new situations. In other words, these features are engines of innovation.

Models of the immune system fall within the range of two main philosophies. One presents immunity as a battle. The immune system detects and seeks to destroy intruders. The second view is that the immune system is an ensemble that seeks to stay in balance. Its response to whatever affects it is a rebalancing that in fact becomes adaptation [3].

In the first view, the immune system is inflexible in its identity. A successful cure leaves the system as it was before the attack, except that now it retains a memory of the aggressor and will be better able to respond to a similar attack in the future. This perspective goes back to the 1950's when immunology became known as the science of the self-nonself discrimination.

Stephanie Forrest and her group at the University of New Mexico developed a system of negative-detection of intruders in computer systems. Algorithms search the computer system for possible changes in specific areas. Then the system tags the change as an intrusion, deals with it, and retains a memory of the event for quicker future response to a similar intrusion. The immediate practical goal of this approach is to give computers a level of self-protection from unexpected viruses and other harmful intrusions. Forrest and Hofmeyr acknowledge that immunology is more than self-nonself discrimination. In a footnote to "Immunology as Information Processing" [4], they note, however, that their approach is an oversimplification. There are harmless intruders that the immune system tolerates, and there are also immune reactions against non-intruding self-cells as in some cancers and autoimmune responses to what should not be perceived as an attack.

The model that Forrest and Hofmeyr followed is non-adaptive. Although the environment can be dynamic and the attacks novel, the negative-selection response is always the same. Nonself is detected and attacked to destroy it. This is a war type model. It hinges on the destruction of the enemy. It is non-adaptive because in the end, there is nothing to adapt to. Nevertheless, the algorithm of negative-selection is considered a milestone in the development of AIS, and rightly so given the enormous complexity of immunity [1].

Tarakanov and Dasgupta used the approach of self-nonself recognition to develop a formal model of an artificial immune system as a highly distributed learning system [6, 1, 7]. They developed a very simplified artificial formal protein that includes torsion factors as a way to model the shape matching that is one tool for identifying non-self intruders. This is their attempt to model the fact that we have tagging proteins that learn by matching shapes with intruders and binding where a match occurs to act as a tag. Tarakanov and Dasgupta then added reproduction and death of the formal proteins to expand their search function. But this mathematical formalism has not yet been tested out. The problem is that it also needs an environment to function and evaluate the performance. In other words, it is not sufficient to formalize the search proteins; the searchable environment has to be also formalized. Or at least a suitable interface has to be developed to interact with the environment. This is still a difficult task to achieve.

Dasgupta and González compared negative selection and positive selection algorithms in the self-nonself approach to intrusion detection in computer networks [8]. They concluded that elimination of what is different is easier to implement and requires fewer resources than checking positive membership in the self-group. The reason is that positive selection is far more memory intensive, at least in the model they developed. The notion of self was not an abstraction but the membership itself. To detect membership in the set, it is necessary to compute the fitness with respect to each member of the self-set and a degree of abnormality is established. There is self-recognition if the abnormality is smaller than a set limit. Negative recognition is done, however, with respect to a single standard, and therefore is easier to implement but if offers less flexibility. It is interesting that in this approach non-self is identified as novelty and eliminated. I will explore how to take advantage of novelty. Rather than eliminate it, it can be used for learning and adaptation to the novelties of a changing environment. In this way a system can innovate.

Melanie Mitchell developed with Douglas Hofstadter an adaptive model based on analogy making [4, 17]. Although it predated the rise of AIS, it has a feature that can be applied to the construction of an immune system, and that is the development of a balance between exploration and attack. She based this on John Holland's concept of a balance between “exploration and exploitation” [5]. She notes that the immune system exploits information from attacks and allocates resources, but it also continues to explore possibilities that might happen by continuously upgrading an enormous repertoire of cells used in recognizing intruders and initiating immune responses. In an actual response, the immune system tinkers with the intruder using detector or tagging cells until gradually the response becomes increasingly focused. Her constructed system uses random combinations and random mutations to perform such exploration and tinkering. This is in fact a blind search. It is feasible in small domains, but entirely random responses simply would take too long for our immune system to work consistently. Mitchell imagines that the next step is the development of more tailored explorations. What is needed is more preliminary interaction between the unknown organisms and the immune system to reach some sort of cooperative balance. After all, the most successful intruders are those that don't destroy the system, because they would also destroy themselves in the process.

At the opposite end of the immunology spectrum is the systemic model best represented by the initial work of Francisco Varela and then modified in collaboration with Antonio Coutinho. Varela did not follow the self-nonself procedure of identification of foreign agents invading the body. Instead he considered the body a semiotically closed system whose task it to maintain itself in balance. In such a system there are no intrusions, only perturbations to its balance. So the immune reaction is not to a foreign agent as such but to a disturbance. The striking difference between this model and the self-nonself one is that it does not follow the war metaphor. Instead it follows an interactive and cooperative approach. The system reacts to the disturbance and affects not just the agent but also the entire system. This requires the tinkering that Mitchell had noted, and the gradual learning of the situation in search for a solution. In this process, it does not matter what happens to the foreign agent as long as the system regains its balance. In other words, the agent may be incorporated into the system somehow, as long as it does not continue to unbalance it. But now the agent itself may be part of the new balance. In this way, the system adapts and evolves while following its task of maintaining a balance.

I can think of two examples to give a sense of how this systemic approach works. One is the population dynamics in U.S.A. We see that immigration rules act as a regulatory system for the country's semi-permeable borders. Immigrants may or may not enter according to rules, and the systemic reaction is not always one of elimination. Instead, the cultural disturbances of immigrants are transformed into the fabric of the country while at the same time the country readapts to the variations. This leads to the gradual series of adaptations that make the U.S.A. a self-similar nation yet at the same time one that has striking cultural differences from what it was before.

A second example is the Michael Crichton's technofiction novel *The Andromeda Strain*. An unknown strain from outer space has infected a small town and with a destructiveness that surpasses a plague. The zone is isolated and authorities try to figure out what to do to contain the threat, learn about it, and in the end eliminate it. Nothing seems to work. Towards the end, the containment perimeter is breached. An aircraft flies too low and undoubtedly carried the strain with it to its landing site. But surprisingly, no damage results. It turns out that the strain has mutated gradually into a form that is harmless to its new environment and is able to coexist on Earth. What Crichton showed is how the strain itself adapted to the new environment in a sort of highly intelligent systemic immunity in reverse. But we gather that this mutation responds to the devastation that the strain caused in its original form. In other words, the strain mutated from a non-cooperative configuration to a cooperative one, and we can imagine that it did not vanish but coevolved and thrived instead quietly in its new form.

3. IMMUNITY, ADAPTATION, AND INNOVATION

Let's consider balanced interactions in which there is possible co-adaptation of foreign agents, internal disruptions, and a system. Here conflict may be enhancing in some ways, and foreign agents, rather than being erased or assimilated, become contributors to the adapted fabric of the dynamic system.

What I would like to explore here is how immunity functions in this situation. This leads to the question of innovation, because the immune system has to deal with situations it did not encounter before and must balance them somehow without destroying itself. The solution would be an innovation for the system since it carried out a new procedure and since it recognizes the procedure as new and can implement it quickly the next time a similar situation happens.

Traditionally, intrusions of foreign agents and internal disruptions are imperfections in the system and called for corrections to restore the original sense of perfection. But now we know that such state of perfection can only be defined in a negative way and this leads to a non-interactive architecture, inflexible, and not suitable for operating in a changing environment.

As we saw before, a positive definition of self or system is redundant, in the sense that it hinges on membership [8]. But this is the type of sense of self we need because it is flexible. I can change in a gradual and somewhat controlled way. Membership hinges on a preset fuzzy distance to what we can imagine as the center of mass of the existing membership set. As new members join the set, the center of mass shifts. This in turn changes membership acceptance. We can say that the system's self adapts to its environment this way.

If we place this dynamic sense of self as a model within the elementary loop of function, or ELF architecture, that Meystel and Albus developed [9], then we have the foundations for what could be the primary controller of an autonomous system capable of operating in a changing environment.

This is where immune concepts come in. Imagine the low resolution, control ELF as a cell with semi-permeable boundary defined by its membership class and the fuzzy measure that allows new members to enter the cell. Constructed immune mechanism must keep such self in balance internally and also with respect of the rest of the autonomous system that the ELF controls. This immune system then has to work within the top ELF, so to speak, as well as throughout the entire autonomous system, to make sure that shifts in self-identity at the top do not wreck the system. This should be achieved through immune mechanisms interacting through loops [10] at two levels: within the top ELF cell and throughout the entire autonomous system. Perhaps there could be a third loop involving the autonomous system and its environment.

In previous works I proposed that imperfections could fuels innovation through compensatory and repair mechanisms that function like an immune system of the imagination [11, 12]. Innovation could then be pictured as stemming from an immune mechanism coupled with a selection process and a well-tuned construction system. What we need to do is not filter out novelties that penetrate the control ELF, but somehow gauge them, and incorporate what works by turning it into a system innovation. As Dasgupta and González [8] noted, this requires more resources, particularly memory, but is also more effective than using the negative approach that eliminates all novelty.

If we try to incorporate immune capabilities into an autonomous system, the previous review suggests that to behave as an adaptive system in a changing environment, it would need special features in its architecture.

1 – It would need at least one control ELF that serves as self. As suggested before, the self-set uses a fuzzy measure for positive recognition of membership. The autonomous system is launched with an initial self-set. As the system interacts with the changing environment, it becomes infected, so to speak. The control ELF takes in new self-elements that fall within the fuzzy measure for positive recognition of membership. This process is Building blocks of elementary functioning loops that have memory and remodeling capacities. Neural nets can produce these capacities. A layered network of such blocks so that there are external and internal inputs at all levels of the system.

2 – To simplify the complexities of immune interactions, the system should use encoding or tagging mechanisms [13]. This helps integrate the many layers, hierarchies of ELFs and other mechanisms of the system. Such tags are like names in a language. They abstract features and allow operating directly on the tags. As Meystel indicates, this process of generalization and representation helps establish links among different levels of resolution in the system's architecture. It also reduces the complexity of the system because it can operate on compressed representations rather than on originals. The original can be decompressed when needed by calling the tag.

3 - Network loops can synchronize the entire system so that it can interact in various ways, especially in the formation of higher-level memories, remembrances, and remodelings [10]. For this we have the neurological model that Edelman calls reentry, which is the synchronous firing of widely dispersed neurons in the brain. Feedforward and feedback signals produce such reentry loop. We also have examples from the immune system. The intrusion into the system of infections puts into motion feedforward and feedback processes throughout the entire immune system until some final state is reached. The immune system acts as a network that interacts with the infection, learns from it, and readjusts. If all goes well, the system learns from the incident and is better prepared to deal with similar ones it might encounter.

4 - Redundancy [13]. This enhances the capacity of subsystems for self-repair or compensation. Meystel describes redundancy as excess in a system. He notes that although such excess may appear as a waste of resources, it is necessary for exploration. He links redundancy with a certain playfulness that operates in the excess. He links it also to a sense of desire or emotion that can serve as a vague guide in exploring by using excess resources.

Solé et al note that redundancy should be understood not as an excess of the same resources but of variant resources [16]. Edelman called this "degeneracy" [19]: structurally different components can yield similar results. Solé indicates that degeneracy is deeply related to tinkering in evolution in the sense that different systems can perform similar functions and therefore can be made available in non-linear ways to yield solutions that open the possibility of divergence of architectures. In a changing environment, the possibility to carry out a similar task using different architectures is a very robust and adaptive feature, since with changes new architectures would be more suitable to the shifting context and yet the original function can still be performed. Although I imagine that the original task would also be adapting. This dynamics is closely related to neutral development and play.

An adaptive autonomous system depends to some degree on being enabled to carry out the following two processes:

1 - Neutral development [14, 15]. This is a stand-by mode of search and development for no immediate use, but one whose task is to prepare the system for future imbalances by giving it more diversified resources.

Motoo Kimura is the architect of the notion of neutral evolution and he applied it initially at the molecular level. He notes that the overwhelming majority of changes in nature are not caused by natural selection but "by random fixation of selectively neutral or nearly neutral mutants" [14]. He adds that "although such random process are slow and insignificant for our ephemeral existence, in the span of geological times, they become colossal."

In constructed autonomous systems it is possible to speed up neutral search when necessary. This would be possible at moments the system is rather inactive and therefore has more resources available. Lobo, Miller, and Fontana picture neutral search as an aimless process that happens within a landscape of solutions that have optimal peaks [15]. Imagine a static environment for an autonomous system. The system has interacted with the environment and provided optimal solutions for the various regions of the landscape. After this optimal adaptation, the system has nothing else to do. It is stuck, so to speak, with the peaks. Should the landscape change, the system would have to start all over again. In neutral mode, the system does not rest even if the landscape remains static. The system roams around the peaks and develops imperfect solutions. Viability rather than optimization is what counts. In this way, the autonomous system engages in neutral development that later can impact its behavior when confronted with unexpected situations, and make it simpler to find new solutions since it has already tested out many. What fuels all this neutral development work at a time of system leisure is play.

2 - Tinkering and play [12, 16, 20, 21]. Solé et al noted that tinkering is an important process of evolution, and therefore of adaptation. They describe tinkering as the re-use of different parts of a system in order to achieve a given function. Those different parts are put to novel uses to have them do together what they were not really designed to do. Tinkering then is an imaginative use of resources, not really in optimal ways, but as back-ups or alternatives. This contributes to the robustness of the system because redundant subsystems evolve through tinkering, and they may come in handy one day when all normal channels fail. Internal tinkering can give a system more protection against random failure. This process is a clear example of neutral search.

Tinkering is a form of play restricted to a specific set of elements at hand. Play is more general. Piaget offered perhaps the most concise model of adaptive play as pure assimilation. That is to say, the player interacts and absorbs aspects of the world without changing its action schemata. The question is how to implement such concept with autonomous systems. Let's divide the process into two components. One is the unchanging schemata that provides equal weight to all choices available. The other is the individual player with preferences that mark individuality and that would prefer certain choices without clear justification beyond personal preference. Combining these two aspects, the simplest technical way to define play that I can think of, is to model it as a random search with preferences.

The practical advantage of having a play mechanism in an autonomous system whose tasks include exploration in a changing environment, is that a narrowed random search can save a great amount of testing and memory. The preferences arise from the interaction of the self with the local environment. It is a situated preference relevant to the moment. This is the engine that drives neutral development as well as recombinant process of exploration. It also helps the system with the overall capacity for self-repair or compensation. The system's degree of play can be controlled through resource allocation given its current state and task. This can be a control mechanism weighed internally. It could also be weighed externally but this reduces the autonomy of the system.

4. CONCLUSION

Natural immune mechanisms are a source of modeling inspiration for the design of autonomous systems capable of operating in dynamic environments. The construction of AIS processes is just starting. Most work is still at the conceptual level. We saw that natural immunity is an intelligent, learning ensemble of interacting mechanisms, and has a strong adaptive function. Modeling immunity can help make adaptive autonomous able to function in changing environments. We noted that to begin to tap this potential the AIS must be designed not in a negative reactive way following the war model, but in a positive way that is more like balanced trade across differences. As Holland noted, an adaptive system has to be able to exploit as well as explore. We saw that immunity is the part of the system that can assist with exploration as well as help maintain the system's integrity so that it can continue to carry out its tasks.

I focused specifically on two processes associated with immunity that could be implemented in autonomous constructions, although their importance is only indirectly related to achieving assigned tasks. These processes are neutral search, tinkering, and play. All are essential parts of natural immune systems, but are difficult to implement in constructed systems, and may seem to be too much idle luxury given the potentially high cost of making them work well.

But in the research and development phase, when discovery outweighs implementation, and when our imagination is freer in the use of resources, we do have more room to play. We can tinker with models and architectures in unexpected ways as long as dogmas don't hold us back and we keep to the task of allowing innovation to happen. In the case of immune features, they might lead to surprising integrations with architectures we already have. Perhaps by enhancing the capacity of constructed systems to engage in neutral tinkering and play we may render them more curious and, paradoxically, more robust. This enactment of interaction may well mark the transition to a new generation of autonomous systems able to explore, innovate, and adapt while carrying out their tasks.

REFERENCES

[1] Llinás, Rodolfo. *I of the Vortex: From Neurons to Self.* Cambridge: MIT Press, 2001.

[2] Tarakanov, Alexander, Victor Skormin and Svetlana Sokolova, *Immunocomputing: Principles and Applications*. New York: Springer-Verlag, 2003.

[3] Coutinho, António, "A Walk with Francisco Varela from first- to second-generation networks: In search of the structure, dynamics and metadynamics of an organism-centered immune system." In *Biol Res* 36: 17-26, 2003.

[4] Segel, Lee A. and Irun R. Cohen. *Design Principles for the Immune System and Other Distributed Autonomous Systems*. New York: Oxford Univ. Press, 2001.

[5] Holland, J. H. *Adaptation in Natural and Artificial Systems*. 2nd ed. Cambridge: MIT Press, 1992.

[6] Tarakanov, Alexander and Dipankar Dasgupta. "A Formal Model of an Artificial Immune System." In *BioSystems* 55 (1-3), 151-158, 2000.

[7] Dasgupta, D. ed. *Artificial Immune Systems and Their Applications*. New York: Springer-Verlag, 1999.

[8] Dasgupta, Dipankar and Fabio González. "An Immunity-Based Technique to Characterize Intrusions in Computer Networks." In *IEEE Transactions of Evolutionary Computation*, 6 (3), 1081-1088, 2002.

[9] Meystel, Alexander M. and James S. Albus. *Intelligent Systems: Architecture, Design, and Control*. New York: John Wiley, 2002.

[10] Arata, Luis O. "Creation by Looping Interactions." *M/C: A Journal of Media and Culture* 5.4 (2002) http://www.media-culture.org.au/0208/creation.html

[11] Arata, Luis O. "Interactive Measures and Innovation." In *Measuring the Performance and Intelligence of Systems: Proceedings of the 2003 PerMIS Workshop*. NIST Special Publication forthcoming in 2004.

[12] Arata, Luis O. "Interaction, Innovation, and Immunity: Enabling Agents to Play." In *American Association for Artificial Intelligence 2004 Spring Symposium Series: Interaction Between Humans and Autonomous Systems over Extended Operation*, 41-46, 2004.

[13] Meystel, Alexander. "Evolution of Intelligent Systems Architectures: What Should Be Measured?" In *Measuring the Performance and Intelligence of Systems: Proceedings of the 2000 PerMIS Workshop*. NIST Special Publication 970, 361-382, 2001.

[14] Kimura, Motoo. *The Neutral Theory of Molecular Evolution*. Cambridge: Cambridge University Press, 1983.

[15] Lobo, José, John H. Miller, and Walter Fontana, "Neutrality in Technological Landscapes." 2004 (soon to be published).

[16] Solé, Ricard, Ramon Ferrer-Cancho, Jose Montoya, and Sergi Valverde, "Selection, Tinkering, and Emergence in Complex Networks." In *Complexity* 8, No. 1, 20-33, 2002.

[17] Mitchell, Melanie. *Analogy-Making as Perception: A Computer Model*. Cambridge: MIT Press, 1993.

[18] Mitchell, Melanie and Stephanie Forrest. "Genertic Algorithms and Artificial Life." In *Artificial Life* 1 (3), 267-289, 1994.

[19] Edelman, Gerald and Giulio Tononi. *A Universe of Consciousness: How Matter Becomes Imagination*. New York: Basic Books, 2000.

[20] "Modeling Interactive Intelligences." In *Measuring the Performance and Intelligence of Systems: Proceedings of the 2002 PerMIS Workshop*. NIST Special Publication 990. September 2002. http://www.isd.mel.nist.gov/research_areas/research_engineering/Performance_Metrics/PerMIS_2002_Proceedings/Arata.pdf

[21] "Can Your Autonomous Robot Come Out and Play?" In *Integration of Knowledge Intensive Multi-Agent Systems. KIMAS 03: Modeling, Exploration, and Engineering*. Institute of Electrical and Electronics Engineers IEEE, 2003.

[22] "Reflections on Interactivity." In *Rethinking Media Change: The Aesthetics of Transition*. David Thorburn and Henry Jenkins, eds. Boston: MIT Press, 2003.